

Review Article

Longevity in the age of artificial intelligence: opportunities, challenges, and the path toward equitable implementation

Nasir Dehghan ^{1*}, Yeganeh Siyar ¹

¹ Nutrition and Health Research Center, Nasim Institute, Tehran, Iran.

*Corresponding author and reprints: Yeganeh Siyar, Nutrition and Health Research Center, Nasim Institute, Tehran, Iran.

Email: siyaryeganeh@gmail.com

Received: 25 Oct 2025

Accepted: 29 Jun 2026

Published: 11 Oct 2026

Abstract

Background: The progressive integration of artificial intelligence (AI) into healthcare systems has generated substantial interest in its potential to extend human healthspan and lifespan. This narrative review examines the current state of AI applications relevant to longevity, with a focus on disease prediction and prevention, drug discovery, genomic analysis, and age-related risk stratification in older adults.

Methods: A comprehensive review of peer-reviewed literature was conducted using PubMed, Scopus, and Web of Science databases, supplemented by gray literature from regulatory agencies and international health organizations. Studies published between 2013 and 2024 were prioritized, with foundational works included where appropriate.

Results: Artificial intelligence demonstrates considerable promise in early detection of cardiovascular disease, oncologic conditions, and neurodegenerative disorders, often outperforming conventional risk models. Machine learning-based biological aging clocks offer novel frameworks for quantifying individual aging trajectories. However, persistent challenges including data fragmentation, algorithmic bias, regulatory ambiguity, and limited clinical interpretability constrain widespread adoption.

Conclusion: While AI holds transformative potential for extending both lifespan and healthspan, its realization depends on rigorous validation, equitable data representation, transparent model design, and interdisciplinary governance frameworks. AI should be positioned as a clinical decision-support tool that augments, rather than supplants, human medical judgment.

Keywords: Aging; Artificial Intelligence; Disease Prevention; Healthcare; Longevity; Predictive Machine Learning.

Cite this article as: Dehghan N, Siyar Y. Longevity in the age of artificial intelligence: opportunities, challenges, and the path toward equitable implementation. *Soc Determinants Health*. 2026;12(1):1-11. DOI: <http://dx.doi.org/10.22037/SDH.v12i1.50696>

Introduction

The aspiration for extended longevity has characterized human civilization across recorded history; however, the convergence of computational technology and biomedical science now renders this ambition scientifically tractable

in unprecedented ways. According to United Nations projections, approximately two billion individuals globally will be aged 60 years or older by 2050, representing more than a twofold increase relative to 2015 figures [1]. This

demographic transformation necessitates a fundamental reorientation of healthcare delivery systems toward proactive, personalized, and scalable interventions.

Contemporary evidence documents remarkable gains in life expectancy over the past century. Where average lifespans in most regions once approximated 40–50 years, global life expectancy has risen to approximately 72.8 years, with high-income nations frequently exceeding 80 years [2]. These gains reflect multifactorial advances, including the management of cardiovascular risk, oncologic therapeutics, and widespread immunization programs [3,4].

Despite these achievements, the burden of non-communicable, chronic diseases—including type 2 diabetes mellitus, cardiovascular disease, cerebrovascular accident sequelae, and musculoskeletal disorders—continues to erode both lifespan and health-related quality of life [5,6]. Reactive healthcare models, which intervene only after pathological manifestations become clinically apparent, are structurally ill-suited to address the complex, multimorbid conditions characteristic of aging populations. Artificial intelligence, by contrast, offers a paradigm shift: the capacity to analyze vast, heterogeneous health datasets, identify subtle pre-symptomatic signals, and generate individualized predictive models [7].

Artificial intelligence encompasses a spectrum of computational methodologies—including supervised machine learning, deep neural networks, natural language processing, and reinforcement learning—that enable systems to identify complex patterns in high-dimensional datasets without explicit rule-based programming. In clinical medicine, machine learning has accelerated the identification of candidate anti-aging compounds [8], elucidated person-specific genomic aging signatures [9], and enabled risk stratification for age-associated

conditions with a precision that substantially exceeds that of traditional actuarial tools [10].

However, the deployment of AI in clinical settings involves significant ethical, regulatory, and operational complexities. Concerns regarding data security, algorithmic fairness, model generalizability, and the digital stratification of healthcare access demand rigorous examination commensurate with the enthusiasm surrounding these technologies [11].

This review synthesizes available evidence on AI's role in promoting human longevity, examining applications in disease prediction, prevention, and biological aging assessment. It critically appraises both the documented achievements and the substantive barriers to implementation, with the aim of providing a balanced, evidence-based account for researchers, clinicians, and policymakers engaged in the emerging field of AI-augmented longevity medicine.

2. Artificial Intelligence in Disease Prediction and Prevention

Among the most consequential applications of AI in longevity medicine is its capacity for early disease detection—identifying pathological risk prior to the emergence of clinical symptoms. Machine learning models, trained on retrospective datasets of electronic health records, imaging data, and genomic profiles, can identify patterns associated with adverse outcomes at a stage when preventive intervention remains most efficacious [12].

2.1 Cardiovascular Disease

Cardiovascular disease remains the leading cause of mortality globally. Classical risk stratification instruments, such as the Framingham Risk Score, rely on a limited number of discrete variables and demonstrate restricted discriminative accuracy across diverse populations. Artificial intelligence models that integrate

genetic, behavioral, longitudinal laboratory, and wearable device data have demonstrated superior performance in identifying individuals at elevated cardiovascular risk [13]. Multiple independent studies have confirmed that machine learning approaches significantly outperform traditional models in predicting major adverse cardiovascular events [14].

2.2 Oncologic Detection

Artificial intelligence is increasingly deployed in oncologic screening, where early detection correlates directly with improved prognosis. Convolutional neural networks trained on dermatoscopic images have achieved diagnostic accuracy for malignant melanoma equivalent to board-certified dermatologists [15]. Analogous systems have demonstrated efficacy in the analysis of mammographic, pulmonary, and colonoscopic imagery for early malignancy detection [16]. These developments suggest that AI-assisted screening could meaningfully advance the temporal boundary of cancer diagnosis toward stages associated with curative therapeutic potential.

2.3 Neurodegenerative Conditions

Neurodegenerative disorders, particularly Alzheimer's disease, frequently progress subclinically for years prior to diagnosis, when intervention is most likely to meaningfully alter disease trajectory. Machine learning algorithms that integrate neuroimaging biomarkers, polygenic risk scores, and neuropsychological assessment data have demonstrated capacity to identify individuals at risk for cognitive decline substantially in advance of symptom onset [17]. Earlier identification creates a therapeutic window that is absent in reactive diagnostic paradigms.

2.4 Metabolic and Cerebrovascular Disease

In the domain of metabolic disease, machine learning models trained on electronic health records have

demonstrated high accuracy in identifying individuals at risk for type 2 diabetes, enabling targeted preventive interventions [18,26]. Similarly, AI-driven stroke prediction models that integrate multimodal clinical variables have outperformed standard risk calculators in prospective validation cohorts, potentially enabling more effective primary prevention [19].

2.5 Infectious Disease Surveillance

The COVID-19 pandemic provided compelling evidence of AI's utility in infectious disease management. Machine learning systems rapidly identified vulnerable subpopulations and optimized resource allocation during peak demand periods [20]. These capabilities extend to other infectious diseases, where AI-assisted analysis of clinical and microbiological data supports timely differential diagnosis and pathogen-specific treatment selection [21].

2.6 Biological Aging Clocks

A particularly innovative application involves computational models that estimate biological age independent of chronological age. Trained on multi-omic datasets—including DNA methylation profiles, transcriptomic data, and clinical biomarkers—these "aging clocks" produce estimates of physiological aging rate that correlate with health outcomes and mortality risk [22,23]. Some investigators propose that such models may ultimately serve not merely as predictive instruments but as dynamic readouts for evaluating the efficacy of candidate geroprotective interventions [23].

3. Artificial Intelligence in Drug Discovery and Genomic Analysis

Artificial intelligence has materially transformed both the rate and scope of pharmaceutical discovery. Conventional drug development pipelines are characterized by extended timelines and high attrition rates; AI-driven virtual screening and generative molecular design

platforms substantially reduce these barriers by rapidly traversing vast chemical spaces to identify promising candidate compounds [8].

In the specific context of aging biology, machine learning systems have been applied to identify compounds that modulate candidate longevity pathways, including mechanistic target of rapamycin (mTOR) signaling, sirtuin activity, and senolytic mechanisms. These *in silico* approaches enable hypothesis generation at a scale and speed that would be unachievable through conventional experimental paradigms.

Deep learning architectures have similarly advanced genomic analysis, enabling systematic identification of non-linear, high-order interactions among genetic variants and their phenotypic correlates [9]. The elucidation of individual-specific genomic aging signatures—patterns of variant accumulation, epigenetic modification, and telomeric attrition—provides a foundation for precision geroprotection tailored to each patient's biological profile. The integration of polygenic risk scoring with environmental and behavioral variables represents an emerging frontier in personalized preventive medicine [39].

4. Challenges in Implementing Artificial Intelligence in Healthcare

Despite considerable promise, the integration of AI into clinical practice is impeded by a constellation of substantive challenges that demand systematic attention.

4.1 Data Quality, Availability, and Representativeness

The performance of any machine learning system is fundamentally contingent on the quality and representativeness of its training data. Electronic health record data are frequently fragmented across disparate institutional repositories, encoded in non-interoperable formats, and subject to

systematic documentation errors and omissions [43]. These limitations constrain the generalizability of models trained on such data.

A more consequential concern is the underrepresentation of specific demographic groups—including racial and ethnic minorities, older adults, and individuals from low-income settings—in the datasets used to develop AI systems [44]. Such representational gaps generate models with differential performance across subpopulations, with the potential to exacerbate existing health disparities rather than ameliorate them.

4.2 Privacy, Ethics, and Algorithmic Bias

Medical data constitute among the most sensitive categories of personal information, and their incorporation into AI systems raises legitimate concerns regarding unauthorized access, secondary use, and re-identification risk [45]. Legislative frameworks, including the Health Insurance Portability and Accountability Act (HIPAA) in the United States and the General Data Protection Regulation (GDPR) in the European Union, establish minimum standards for data governance, but the tension between privacy preservation and the data volume requirements of AI training remains unresolved.

Beyond privacy, AI systems raise substantive ethical questions regarding accountability for diagnostic or therapeutic errors attributable to algorithmic outputs, the propagation of historical inequities embedded in training data, and the adequacy of informed consent frameworks for AI-assisted clinical decisions [11,46].

Empirical evidence of algorithmic bias in clinical AI is documented and concerning. A landmark analysis by Obermeyer et al. demonstrated that a widely deployed commercial risk stratification algorithm systematically underestimated the medical need of Black patients relative to White patients of equivalent disease severity,

because it employed healthcare expenditure—a proxy confounded by systemic access inequities—as a surrogate for clinical complexity [47]. This finding illustrates that algorithmic bias in healthcare is not merely a theoretical concern but an active mechanism capable of amplifying entrenched disparities.

4.3 Regulatory Frameworks and Clinical Validation

The regulatory evaluation of AI-based clinical decision-support tools presents methodological challenges that current frameworks, largely designed for static pharmaceutical agents and fixed-function medical devices, do not fully address [48]. Adaptive, continuously learning AI systems may exhibit performance drift over time, necessitating longitudinal post-deployment monitoring protocols that extend beyond the scope of traditional clinical trial designs.

A further concern is the performance gap between controlled research settings and real-world clinical environments. Models that demonstrate strong performance in retrospective validation may fail to generalize when confronted with the heterogeneity of clinical populations, documentation practices, and institutional workflows [49]. Multi-site prospective validation represents the methodological standard required before clinical deployment, yet few published AI systems have undergone evaluation of this rigor.

4.4 Clinical Integration and Physician Acceptance

Even technically validated AI tools may encounter substantial barriers to clinical adoption. Physicians trained prior to the widespread deployment of computational decision-support tools often regard algorithmic recommendations with justifiable skepticism, particularly when the reasoning underlying these recommendations is opaque [50]. This "black box" problem—characteristic of deep learning architectures with millions of

learned parameters—impedes the development of physician trust necessary for effective clinical integration [51].

Research on explainable AI (XAI) seeks to address this limitation by developing methods that render model reasoning interpretable to clinical end-users. However, the trade-off between predictive performance and interpretability remains an active area of methodological inquiry, and no consensus approach has yet emerged [52]. Successful clinical integration additionally requires that AI tools be embedded within existing workflows in ways that reduce rather than augment cognitive and administrative burden [53].

4.5 Economic Considerations and Health Equity

The development and maintenance of high-performance AI systems requires substantial investments in computational infrastructure, data engineering, and specialized human capital [54]. Smaller institutions and health systems in resource-constrained settings risk systematic exclusion from the benefits of these technologies, potentially concentrating AI-mediated health gains among already-advantaged populations. Furthermore, AI models require ongoing retraining as patient populations, practice patterns, and disease epidemiology evolve, necessitating sustained organizational and financial commitments [55].

4.6 Interpretability, Accountability, and Decision Support

The interpretability challenge in clinical AI extends beyond physician trust to encompass questions of accountability and continuous quality improvement. When an AI system contributes to a diagnostic error or adverse clinical outcome, the absence of transparent reasoning impedes root cause analysis and corrective action [56]. Some scholars have argued that intrinsically interpretable models should be preferred for high-stakes clinical decisions, accepting a

modest reduction in predictive accuracy in exchange for transparency [51].

5. The Path Forward: Strategic Recommendations

Addressing the foregoing challenges requires coordinated action across the research, clinical, regulatory, and policy domains. The following recommendations are advanced on the basis of the reviewed evidence.

5.1 For Researchers and AI Developers

Research teams should adopt a design philosophy that prioritizes clinical transparency alongside predictive accuracy, particularly for applications involving high-stakes decision-making. Prospective identification and mitigation of algorithmic bias across demographic subgroups should be treated as a methodological requirement rather than an optional extension. Clinical validation conducted across diverse institutional and geographic settings should be considered a prerequisite for claims of generalizability. Adoption of open-science principles—including code and dataset sharing where privacy constraints permit—would accelerate external replication and strengthen the evidentiary base for clinical deployment.

5.2 For Clinicians and Healthcare Institutions

Healthcare institutions should invest in structured AI literacy training that equips clinicians with the conceptual tools necessary to critically evaluate algorithmic recommendations and recognize the limitations of specific systems. AI should be implemented as a component of existing clinical workflows rather than as a supplementary layer that increases documentation burden. Systematic post-deployment performance monitoring, including surveillance for performance drift and differential outcomes across patient subgroups, should be operationalized as a condition of sustained use.

5.3 For Policymakers and Regulators

Regulatory agencies should develop adaptive frameworks capable of evaluating continuously learning AI systems and applying post-market surveillance requirements appropriate to their dynamic nature. Data governance policies should facilitate the creation of interoperable, privacy-preserving health data infrastructure to support AI development without compromising individual rights. Incentive structures should be designed to promote AI development for underserved populations and resource-limited settings, where the potential for disproportionate benefit is substantial [60,61]. Clear legal and professional frameworks for accountability in cases of AI-attributable harm should be established prospectively.

5.4 For Healthcare Consumers

Individuals should be supported in developing health literacy that encompasses the role of AI in clinical decision-making, including the right to disclosure when AI systems have been used in their care. Participation in AI research trials constitutes a mechanism by which patients can contribute to improving the generalizability and equity of AI systems. Engagement with preventive health recommendations generated by AI-assisted risk assessment represents an opportunity for individuals to realize potential longevity benefits through behavior modification prior to the emergence of clinical disease.

6. Research Priorities

Several domains require prioritized investigation to advance the field.

First, the development and validation of biological aging biomarkers that are both clinically interpretable and responsive to intervention should be considered a central priority. Aging clocks that merely describe chronological deviation provide limited actionable guidance; instruments capable of documenting biological rejuvenation in response to geroprotective interventions

would substantially advance translational aging research [22,23].

Second, multimorbidity—the coexistence of multiple chronic conditions—is the normative experience of older adults and represents a domain where AI's capacity to integrate heterogeneous data streams may offer particular value. AI models designed explicitly for multimorbid clinical contexts, rather than optimized for single-disease performance, warrant dedicated research investment [29].

Third, personalized intervention optimization—the use of reinforcement learning and contextual bandit algorithms to identify the behavioral, pharmacological, and therapeutic strategies most likely to benefit a given individual based on their unique biological and environmental profile—represents a theoretically compelling but empirically underdeveloped application.

Fourth, the development of low-cost, deployable AI tools specifically designed for low- and middle-income country contexts—where the burden of preventable chronic disease is greatest and clinical human capital most constrained—represents both a scientific and an ethical priority [61].

Fifth, rigorous health equity research that examines whether AI-mediated interventions narrow or widen existing health disparities is essential to ensure that the field's aspirational goals translate into equitably distributed outcomes.

7. Ethical Framework

All applications of AI in healthcare must be evaluated against a consistent ethical framework grounded in established principles of biomedical ethics. Beneficence requires that AI tools be deployed with demonstrable intent to improve patient-centered outcomes, rather than primarily to reduce costs or enhance institutional efficiency. Non-maleficence demands that prospective harms—

including bias-mediated disparities and privacy breaches—be systematically identified and mitigated prior to deployment. Respect for patient autonomy requires that individuals be informed of AI's role in their care and retain ultimate decision-making authority over treatment choices. Justice mandates that the benefits and risks of AI be distributed equitably across populations, without discrimination on the basis of socioeconomic status, geographic location, race, or other protected characteristics. Finally, transparency requires that AI systems be designed to the maximum feasible degree of interpretability, enabling meaningful oversight by clinicians, patients, and regulatory bodies.

8. Conclusion

Artificial intelligence represents a scientifically credible and potentially transformative contributor to the extension of human lifespan and healthspan. Its capacity to analyze high-dimensional, heterogeneous health data and generate actionable predictive insights positions it as a powerful complement to conventional clinical expertise across domains including cardiovascular risk stratification, oncologic screening, neurodegenerative disease detection, and personalized pharmacology. Biological aging clocks and AI-assisted drug discovery further suggest the possibility of eventually targeting the aging process itself, not merely its downstream pathological manifestations.

The historical precedents for population-level longevity gains—from the control of infectious disease in the twentieth century to the development of evidence-based cardiovascular preventive medicine—demonstrate that systemic, technology-enabled interventions can produce substantial and lasting extensions of average lifespan. Artificial intelligence may constitute the next such inflection point.

However, the responsible realization of this potential requires confronting substantive

challenges with intellectual honesty and institutional commitment. Data fragmentation, algorithmic bias, interpretive opacity, regulatory inadequacy, and inequitable access are not peripheral concerns to be managed after deployment; they are central conditions for responsible development. AI must be validated with the same methodological rigor applied to pharmaceutical agents, deployed with appropriate governance structures, and monitored continuously for both performance and equity outcomes.

Critically, AI must be understood as a clinical decision-support technology that augments the judgment, compassion, and contextual sensitivity of skilled clinicians—not as a replacement for the human dimensions of medical care. The physician-patient relationship, grounded in trust, communication, and individualized assessment, cannot be computationally replicated and must remain at the center of clinical practice.

The path forward is one of genuine interdisciplinary collaboration: engineers and clinicians, researchers and regulators, health systems and communities. Progress requires investment not only in algorithms and infrastructure but in equitable data ecosystems, adaptive regulatory frameworks, AI-literate clinical workforces, and governance structures that hold technology accountable to the populations it purports to serve. If these conditions are met, artificial intelligence in healthcare may fulfill its extraordinary potential—not as a technological spectacle, but as a measured and equitable contribution to one of medicine's most fundamental aspirations: a longer, healthier life for all.

Acknowledgments

The authors gratefully acknowledge the institutional support of the Education Development Centre at Iran University of Medical Sciences. The authors also recognize the foundational contributions of

the global research community to the evidence base reviewed herein.

Conflict of Interest

The authors declare no conflicts of interest, financial or otherwise, relevant to the content of this manuscript.

Funding

This research received no specific funding from public, commercial, or nonprofit funding agencies.

Authors' Contributions

N.D.: conceptualization, original draft preparation, and overall supervision. Y.S.: literature review, critical appraisal, and manuscript revision. Both authors reviewed and approved the final manuscript.

AI Disclosure

In the preparation of this manuscript, the authors used artificial intelligence language tools as writing assistance. Specifically, Claude (Anthropic, model claude-sonnet-4-20250514, 2025) and ChatGPT (OpenAI, GPT-4, 2023) were employed during initial drafting and text refinement phases. All AI-generated content was critically reviewed, substantially revised, and verified for scientific accuracy by the human authors, who assume full responsibility for the integrity and originality of the final manuscript. This disclosure is made in accordance with emerging guidelines from the Iranian Ministry of Health and Education regarding transparent reporting of AI tool use in academic publications.

References

1. Anthropic. Claude (claude-sonnet-4-20250514) [Large language model]. Anthropic PBC; 2025. Available from: <https://www.anthropic.com/claude>
2. OpenAI. ChatGPT (GPT-4) [Large language model]. OpenAI LP; 2023. Available from: <https://openai.com/chatgpt>
3. United Nations, Department of Economic and Social Affairs, Population Division. World Population Ageing 2020: Highlights. New York: United Nations; 2020.

4. World Bank Group. Life expectancy at birth, total (years) [Internet]. Washington, DC: World Bank; 2023 [cited 2024]. Available from: <https://data.worldbank.org>
5. Roth GA, Mensah GA, Johnson CO, Addolorato G, Ammirati E, Baddour LM, et al. Global burden of cardiovascular diseases and risk factors, 1990–2019: update from the GBD 2019 Study. *J Am Coll Cardiol*. 2020;76(25):2982–3021.
6. Sung H, Ferlay J, Siegel RL, Laversanne M, Soerjomataram I, Jemal A, et al. Global cancer statistics 2020: GLOBOCAN estimates of incidence and mortality worldwide for 36 cancers in 185 countries. *CA Cancer J Clin*. 2021;71(3):209–249.
7. Saeedi P, Petersohn I, Salpea P, Malanda B, Karuranga S, Unwin N, et al. Global and regional diabetes prevalence estimates for 2019 and projections for 2030 and 2045: results from the International Diabetes Federation Diabetes Atlas, 9th edition. *Diabetes Res Clin Pract*. 2019;157:107843.
8. Feigin VL, Stark BA, Johnson CO, Roth GA, Bisignano C, Abady GG, et al. Global, regional, and national burden of stroke and its risk factors, 1990–2019: a systematic analysis for the Global Burden of Disease Study 2019. *Lancet Neurol*. 2021;20(10):795–820.
9. Topol EJ. High-performance medicine: the convergence of human and artificial intelligence. *Nat Med*. 2019;25(1):44–56.
10. Fleming N. How artificial intelligence is changing drug discovery. *Nature*. 2018;557(7706):S55–S57.
11. Eraslan G, Avsec Ž, Gagneur J, Theis FJ. Deep learning: new computational modelling techniques for genomics. *Nat Rev Genet*. 2019;20(7):389–403.
12. Beam AL, Kohane IS. Big data and machine learning in health care. *JAMA*. 2018;319(13):1317–1318.
13. Char DS, Shah NH, Magnus D. Implementing machine learning in health care — addressing ethical challenges. *N Engl J Med*. 2018;378(11):981–983.
14. Obermeyer Z, Emanuel EJ. Predicting the future — big data, machine learning, and clinical medicine. *N Engl J Med*. 2016;375(13):1216–1219.
15. Krittanawong C, Zhang H, Wang Z, Aydar M, Kitai T. Artificial intelligence in precision cardiovascular medicine. *J Am Coll Cardiol*. 2017;69(21):2657–2664.
16. Ambale-Venkatesh B, Yang X, Wu CO, Liu K, Hundley WG, McClelland R, et al. Cardiovascular event prediction by machine learning: the Multi-Ethnic Study of Atherosclerosis. *Circ Res*. 2017;121(9):1092–1101.
17. Esteva A, Kuprel B, Novoa RA, Ko J, Swetter SM, Blau HM, et al. Dermatologist-level classification of skin cancer with deep neural networks. *Nature*. 2017;542(7639):115–118.
18. Bi WL, Hosny A, Schabath MB, Giger ML, Birkbak NJ, Mehrtash A, et al. Artificial intelligence in cancer imaging: clinical challenges and applications. *CA Cancer J Clin*. 2019;69(2):127–157.
19. Rathore S, Habes M, Iftikhar MA, Shacklett A, Davatzikos C. A review on neuroimaging-based classification studies and associated feature extraction methods for Alzheimer's disease and its prodromal stages. *Neuroimage*. 2017;155:530–548.
20. Kavakiotis I, Tsave O, Salifoglou A, Maglaveras N, Vlahavas I, Chouvarda I. Machine learning and data mining methods in diabetes research. *Comput Struct Biotechnol J*. 2017;15:104–116.
21. Douiri A, McKeivitt C, Emmett ES, Rudd AG, Wolfe CD. Long-term effects of secondary preventions on the burden of stroke: the South London Stroke Register. *Circulation*. 2013;128(12):1341–1348.
22. Vaid A, Somani S, Russak AJ, De Freitas JK, Chaudhry FF, Paranjpe I, et al. Machine learning to predict mortality and critical events in a cohort of patients with COVID-19 in New York City: model development and validation. *J Med Internet Res*. 2020;22(11):e24018.
23. Stafford IS, Kellermann M, Mossotto E, Beattie RM, MacArthur BD, Ennis S. A systematic review of the applications of artificial intelligence and machine learning in autoimmune diseases. *NPJ Digit Med*. 2020;3:30.
24. Pyrkov TV, Slipensky K, Barg M, Kondrashin A, Zhurov B, Zenin A, et al. Extracting biological age from biomedical data via deep learning: too much of a good thing? *Sci Rep*. 2018;8(1):5210.
25. Puetz E, Mamoshina P, Aliper A, Korzinkin M, Moskalev A, Kolosov A, et al. Deep biomarkers of human aging: application of deep neural networks to biomarker development. *Aging (Albany NY)*. 2016;8(5):1021–1033.
26. Murdoch TB, Detsky AS. The inevitable application of big data to health care. *JAMA*. 2013;309(13):1351–1352.
27. Goldstein BA, Navar AM, Pencina MJ, Ioannidis JPA. Opportunities and challenges in developing risk prediction models with electronic health records data: a systematic review. *J Am Med Inform Assoc*. 2017;24(1):198–208.
28. Zheng T, Xie W, Xu L, He X, Zhang Y, You M, et al. A machine learning-based framework to identify type 2 diabetes through electronic health records. *Int J Med Inform*. 2017;97:120–127.

29. Weng SF, Reys J, Kai J, Garibaldi JM, Qureshi N. Can machine-learning improve cardiovascular risk prediction using routine clinical data? *PLoS One*. 2017;12(4):e0174944.
30. Qiu S, Joshi PS, Miller MI, Xue C, Zhou X, Karjadi C, et al. Development and validation of an interpretable deep learning framework for Alzheimer's disease classification. *Brain*. 2020;143(6):1920–1933.
31. Miotto R, Wang F, Wang S, Jiang X, Dudley JT. Deep learning for healthcare: review, opportunities and challenges. *Brief Bioinform*. 2018;19(6):1236–1246.
32. Gulshan V, Peng L, Coram M, Stumpe MC, Wu D, Narayanaswamy A, et al. Development and validation of a deep learning algorithm for detection of diabetic retinopathy in retinal fundus photographs. *JAMA*. 2016;316(22):2402–2410.
33. Ting DSW, Cheung CY, Lim G, Tan GSW, Quang ND, Gan A, et al. Development and validation of a deep learning system for diabetic retinopathy and related eye diseases using retinal images from multiethnic populations with diabetes. *JAMA*. 2017;318(22):2211–2223.
34. Graham S, Depp C, Lee EE, Nebeker C, Tu X, Kim HC, et al. Artificial intelligence for mental health and mental illnesses: an overview. *Curr Psychiatry Rep*. 2019;21(11):116.
35. Shatte ABR, Hutchinson DM, Teague SJ. Machine learning in mental health: a scoping review of methods and applications. *Psychol Med*. 2019;49(9):1426–1448.
36. Fitzpatrick KK, Darcy A, Vierhile M. Delivering cognitive behavior therapy to young adults with symptoms of depression and anxiety using a fully automated conversational agent (Woebot): a randomized controlled trial. *JMIR Ment Health*. 2017;4(2):e19.
37. Kuo CC, Chang CM, Liu KT, Lin WK, Chiang HY, Chung CW, et al. Automation of the kidney function prediction and classification through ultrasound-based kidney imaging using deep learning. *NPJ Digit Med*. 2019;2:29.
38. Makino M, Yoshimoto R, Ono M, Itoko T, Katsuki T, Koseki A, et al. Artificial intelligence predicts the progression of diabetic kidney disease using big data machine learning. *Sci Rep*. 2019;9(1):11862.
39. Liu F, Zhou Z, Jang H, Samsonov A, Zhao G, Kijowski R. Deep convolutional neural network and 3D deformable approach for tissue segmentation in musculoskeletal magnetic resonance imaging. *Magn Reson Med*. 2018;79(4):2379–2391.
40. Bien N, Rajpurkar P, Ball RL, Irvin J, Park A, Jones E, et al. Deep-learning-assisted diagnosis for knee magnetic resonance imaging: development and retrospective validation of MRNet. *PLoS Med*. 2018;15(11):e1002699.
41. Das S, Forer L, Schönherr S, Sidore C, Locke AE, Kwong A, et al. Next-generation genotype imputation service and methods. *Nat Genet*. 2016;48(10):1284–1287.
42. Codella NCF, Gutman D, Celebi ME, Helba B, Marchetti MA, Dusza SW, et al. Skin lesion analysis toward melanoma detection: ISIC 2017 challenge. *Proceedings of ISBI 2018, IEEE 15th International Symposium on Biomedical Imaging*. 2018:168–172.
43. Hosny A, Parmar C, Quackenbush J, Schwartz LH, Aerts HJWL. Artificial intelligence in radiology. *Nat Rev Cancer*. 2018;18(8):500–510.
44. Campanella G, Hanna MG, Geneslaw L, Mirafior A, Silva VWK, Busam KJ, et al. Clinical-grade computational pathology using weakly supervised deep learning on whole slide images. *Nat Med*. 2019;25(8):1301–1309.
45. Kelly CJ, Karthikesalingam A, Suleyman M, Corrado G, King D. Key challenges for delivering clinical impact with artificial intelligence. *BMC Med*. 2019;17:195.
46. Gianfrancesco MA, Tamang S, Yazdany J, Schmajuk G. Potential biases in machine learning algorithms using electronic health record data. *JAMA Intern Med*. 2018;178(11):1544–1547.
47. Price WN 2nd, Cohen IG. Privacy in the age of medical big data. *Nat Med*. 2019;25(1):37–43.
48. Morley J, Machado CCV, Burr C, Cows J, Joshi I, Taddeo M, et al. The ethics of AI in health care: a mapping review. *Soc Sci Med*. 2020;260:113172.
49. Obermeyer Z, Powers B, Vogeli C, Mullainathan S. Dissecting racial bias in an algorithm used to manage the health of populations. *Science*. 2019;366(6464):447–453.
50. Gerke S, Minssen T, Cohen G. Ethical and legal challenges of artificial intelligence-driven healthcare. In: *Artificial Intelligence in Healthcare*. Cambridge: Academic Press; 2020. p. 295–336.
51. Liu X, Rivera SC, Moher D, Calvert MJ, Denniston AK; SPIRIT-AI and CONSORT-AI Working Group. Reporting guidelines for clinical trial reports for interventions involving artificial intelligence: the CONSORT-AI extension. *Nat Med*. 2020;26(9):1364–1374.
52. Blease C, Kaptchuk TJ, Bernstein MH, Mandl KD, Halamka JD, DesRoches CM. Artificial intelligence and the future of primary care: exploratory qualitative study of UK general practitioners' views. *J Med Internet Res*. 2019;21(3):e12802.

53. Rudin C. Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead. *Nat Mach Intell.* 2019;1(5):206–215.
54. Holzinger A, Biemann C, Pattichis CS, Kell DB. What do we need to build explainable AI systems for the medical domain? *arXiv preprint arXiv:1712.09923.* 2017.
55. Cresswell K, Bates DW, Sheikh A. Ten key considerations for the successful implementation and adoption of large-scale health information technology. *J Am Med Inform Assoc.* 2013;20(e1):e9–e13.
56. Matheny ME, Whicher D, Thadaney Israni S. Artificial intelligence in health care: a report from the National Academy of Medicine. *JAMA.* 2020;323(6):509–510.
57. Reddy S, Allan S, Coghlan S, Cooper P. A governance model for the application of AI in health care. *J Am Med Inform Assoc.* 2020;27(3):491–497.
58. Cabitza F, Rasoini R, Gensini GF. Unintended consequences of machine learning in medicine. *JAMA.* 2017;318(6):517–518.
59. Chen JH, Asch SM. Machine learning and prediction in medicine — beyond the peak of inflated expectations. *N Engl J Med.* 2017;376(26):2507–2509.
60. Sendak MP, Gao M, Brajer N, Balu S. Presenting machine learning model information to clinical end users with model facts labels. *NPJ Digit Med.* 2020;3:41.
61. Paranjape K, Schinkel M, Panday RN, Car J, Nanayakkara P. Introducing artificial intelligence training in medical education. *JMIR Med Educ.* 2019;5(2):e16048.
62. Chen IY, Pierson E, Rose S, Joshi S, Ferryman K, Ghassemi M. Ethical machine learning in healthcare. *Annu Rev Biomed Data Sci.* 2021;4:123–144.
63. Wahl B, Cossy-Gantner A, Germann S, Schwalbe NR. Artificial intelligence (AI) and global health: how can AI contribute to health in resource-poor settings? *BMJ Glob Health.* 2018;3(4):e000798.